

HETEROGENEOUS DISK STORAGE MANAGEMENT TECHNIQUE

5 CROSS REFERENCE TO RELATED APPLICATION

This application claims priority under 35 U.S.C. 119(e) to U.S. Provisional Patent Application Serial No. 60/392,779, filed July 1, 2002, the teachings of which are incorporated herein.

10

TECHNICAL FIELD

This invention relates to a technique for managing one or more storage devices.

15 BACKGROUND ART

Advances in the design and manufacture of magnetic storage devices have significantly reduced the cost of such devices, as measured in terms of the cost per gigabyte of storage capacity. The relatively low cost of high storage capacity devices has resulted in the proliferation 20 of large arrays of such devices, such as a Redundant Array of Independent Disks (RAID). Various manufacturers of storage devices commonly offer such RAIDs with proprietary software for controlling the individual storage devices within the RAID.

Storage devices from different manufacturers often have different features. Combining storage devices from different manufacturer affords the ability to take advantage of the best 25 features of each device. Unfortunately, no technique presently exists for low-level management of heterogeneous storage devices, such as a set of heterogeneous storage devices within a RAID. While software presently exists for managing RAIDs from different vendors across a network, such software lacks the ability to manage one or more storage devices from separate vendors in a single RAID, or to effectively manage a stand-alone storage device.

30 Thus, there is a need for a technique for managing one or more storage devices.

BRIEF SUMMARY OF THE INVENTION

Briefly, in accordance with present principles, there is provided a method for managing at least one storage device, such as a magnetic disk drive, magneto-optical device, or a tape

5 storage device, and preferably, a plurality of such storage devices. The method commences by identifying each device, typically by its manufacturer. For each identified storage device, a database is created containing information about that device. For example, the information within such a database will include the working characteristics of that storage device, specific interface protocols and operational rules. From the information about each storage device
10 obtained from the corresponding database, a Graphical User Interface (GUI) is created for display to provide a user with at least one menu option for selection. Upon selecting the one menu option, the user obtains at least one of (a) a display of information associated with the identified storage device, and (b) execution of at least one process to control, at least in part, the operation of the identified storage device. The user's selection is processed and the Graphical
15 User Interface is automatically updated in response to the processing of the user's selected menu option. Thus, for example, if the user had selected a menu option to cause the identified storage device to perform a particular operation, the Graphical User Interface provides an updated display that reflects the results of that operation performed by the storage device.

20 BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 depicts a block schematic diagram of a system in accordance with present principles for managing at least one storage device;

25 FIGURE 2 depicts in flow chart form the steps of a method executed by the system of FIG. 1 to manage the at least one storage device; and

FIGURE 3 depicts an opening screen of a Graphical User Interface displayed to a user.

DETAILED DESCRIPTION

30 FIGURE 1 depicts a block schematic diagram of a host system 10 for managing at least one, and preferably a plurality of storage devices, exemplified by storage devices 12₁, 12₂ and 12₃. Each of the storage devices 12₁, 12₂ and 12₃ can take the form of a single magnetic or optical disk drive, a magnetic tape drive, or an array of individual storage devices, such as a

Redundant Array of Independent Disks (RAID). In practice, the storage devices 12₁-12₃ originate from different manufacturers. Thus, each storage device often possesses slightly different operating characteristics that must be taken into account when managing that device.

A fibre channel fabric 14 couples the storage devices 12₁-12₃ to a host bus adapter 16 that 5 interfaces the devices to the host system 10. The host bus adapter 16 typically has its own Small Computer System Interface (SCSI) to funnel information between the host system 10 and the storage devices 12₁-12₃ through the fibre channel fabric 14. The host bus adapter 16 can include other interfaces for disk storage devices that do not use SCSI. Note that the storage devices 12₁-12₃ could each enjoy a direct connection with the host bus adapter 16, thereby obviating the need 10 for the fibre channel fabric 14. Thus, the storage devices 12₁-12₃ are not limited to fibre channel storage devices but could comprise any type of storage device responsive to SCSI commands.

The host system 10 typically comprises a general-purpose computer such as a well known personal computer or a mini computer. Associated with the host system 10 is at least one database (DBS) 18 that contains one or more libraries. Each library within the database 18 15 contains information specific to a corresponding one of the storage devices 12₁-12₃. In particular, the library associated with each of the storage devices 12₁-12₃ typically includes the operating characteristics associated with that storage device. The device operating characteristics includes at least one of (a) operational rules, (b) commands, and (c) processing routines. While the embodiment of FIG. 1 depicts a single database 18 that contains the associated library for 20 each of the storage devices 12₁-12₃, the library associated with each storage device could reside on a separate database (not shown). In practice, each library in the database 18 depends only on SCSI commands. Thus, the host bus adapter 16 need not be limited to a Fibre Channel (FC)/SCSI host bus adapter because regardless of the lower level protocol (e.g., fibre channel fabric in this case), the commands in the library don't need to rely on the specific host bus 25 adapter protocol.

The host system 10 includes a Communication Object Model (COM)/ Distributed Communications Object Model (DCOM) layer 20, which, in practice, takes the form of a memory that includes specific configuration rules for each of the storage devices 12₁-12₃. In practice, the configuration rules for each storage device are determined from the information 30 stored in the corresponding library for each such device. The COM/DCOM layer 20 also includes one or more applets, programs and/or sub-routines (collectively referred as "programs") typically in the form of C++ objects that are exposed via multiple interfaces and can be instantiated via multiple languages across multiple machines in one or more networks. The

programs in the COM/DCOM layer 20 monitor each requested operation for each of the storage devices 12₁-12₃ to ensure that the requested operation complies with the rules for that storage device as established by its manufacturer. Any requested operation of a storage device not in compliance with the rule(s) associated with that device will be blocked and an error message will 5 be generated. Execution by the host system 10 of the programs resident in the COM/DCOM layer 20 can occur without the need for the user to remain at a terminal (not shown).

The host system 10 also executes a Graphical User Interface (GUI) application 22 that manages input information received from each of the storage devices 12₁-12₃ as well as input information received from one or more users. From such input information, the GUI application 10 22 generates a graphical interface for display to the user on a display device (not shown). The overall format of the information displayed by the GUI 22 to the user for each storage drive has a uniform appearance, but different menu options are tailored for each storage device to account for its particular features.

FIGURE 2 illustrates in flow chart form the steps of a method executed by the host 15 system 10 to manage the storage devices 12₁-12₃ in accordance with the present principles. The method of FIG. 2 commences upon execution of step 100 during which the host system 10 of FIG. 1 undertakes discovery of which storage devices are attached thereto. In other words, during step 100, the host system 10 launches a query to determine the identity of the attached devices 12₁-12₃, and particularly, the make (i.e., the manufacturer) and the model number of each 20 device. In addition to discovering the make and model of each device, the host system 10 also ascertains during step 100 whether each device is directly attached to the host bus adapter 16 or is attached through the fibre channel fabric 14. During step 100, the host system 10 will assign an identifier to each discovered storage devices. The assigned identifier can comprise a fibre channel fabric identifier assigned when the device is attached to the fibre channel fabric 14 of 25 FIG. 1. Alternatively, the identifier could comprise a Logical Unit Number (LUN) or vendor identification number assigned to the device.

Following step 100, execution of step 110 occurs at which time, the host system 10 categorizes each discovered storage device by its make and model. During step 110, the host system 10 establishes a library for each storage device discovered during step 100, with each 30 library residing in a single database (e.g., database 18 of FIG. 1) or in separate databases. As discussed previously, each library associated with a storage device will include the operational rules, commands and processing routines associated with that device. Execution of step 120

follows step 110 at which time, the host system 10 executes the Graphical User Interface (GUI) application 22 of FIG. 1 to create and display the GUI to the user.

FIGURE 3 depicts an illustrative screen display 200 that comprises part of the GUI provided to the user. The screen display 200 includes system pane 202 typically appearing at the left-hand side of the screen display. The system pane 202 provides a hierarchical display of the various devices and sub-elements found by the host system 10 during step 100 of FIG. 2, with each device and sub-element appearing in the system pane 202 of FIG. 3 as a corresponding icon. In the exemplary embodiment of FIG. 3, the icon 206 appearing in the system pane 202 corresponds to the presence of a Vendor Unique Storage Device (not shown) found by the host system 10 during step 100 of FIG. 2. The icon 208 appearing in the system pane 202 of FIG. 3 corresponds to a controller node (not shown) within the Vendor Unique Storage Device represented by the icon 206. The icon 210 within the system pane 202 signifies the existence of one or more bound LUNs (not shown) associated with the controller node represented by icon 206. The icon 210 identifies a corresponding LUN within the Bound LUNs represented by the icon 210. Lastly, the icon 212 within the system pane 202 of FIG. 3 signifies the presence of one or more unbound disks (not shown). One or more of the icons 206-212 can have an associated drop-down menus to allow a user to obtain one of (a) a display of information associated with the identified storage device, and (b) execution of at least one process to control, at least in part, the operation of the identified storage device.

Still referring to FIG. 3, the screen display 200 includes a sub-system display pane 214 in addition to the display pane 202. In practice, the sub-system display pane 214 displays specific information related to the selection by a user (via a mouse or the like) of a particular one of the icons 206-212. In the exemplary embodiment, highlighting the controller node icon 208 causes the sub-system display pane 214 to display more specific information about the controller node.

Referring to FIG. 2, following creation and display of the GUI during step 120, the host system 10 then awaits input from a user who typically selects one or more menu options. Such menu options can include (but are not limited to) the following:

- bind a group of physical storage devices into a Logical Unit Number (LUN)
- unbind a set of LUNs
- designate one or more storage devices as hot spares
- set each storage device system clock
- load vendor specific firmware to individual controllers and/or storage devices
- check/set vendor specific controller and/or storage device settings

-6-

- blink the storage device Light Emitting Diodes (LEDs) if present for identification purposes
- identify state information
- establish fibre channel fabric IDs for each storage device
- 5 • start or stop a RAID rebuild operation
- start or stop one or more consistency check operations
- initialize a file system in a discrete storage system

After receipt of by a user of one or more menu selections, then the step 124 of FIG 2 occurs, at which time the host system 10 of FIG. 1 processes the received menu selection. In 10 practice, the host system 10 processes the menu selection through an appropriate software stack until the specific storage device to which the menu selection was directed executes the selection. Following step 124, step 126 occurs, whereupon the host system 10 automatically updates the GUI previously displayed to the user to reflect changes that occurred as a result of execution of the user-selected menu option. In this way, the user can visually observe the results of the 15 selected menu option. Following step 126, program execution branches back to step 122 to wait for user selection of another menu option.

The foregoing describes a technique for managing one or more storage devices having different operating characteristics.